

PCT/IB05/50399

REC'D 03 MAR 2005

WIPO

PCT

PA 1278465

THE UNITED STATES OF AMERICA

TO ALL TO WHOM THESE PRESENTS SHALL COME:

UNITED STATES DEPARTMENT OF COMMERCE

United States Patent and Trademark Office

January 31, 2005

**THIS IS TO CERTIFY THAT ANNEXED HERETO IS A TRUE COPY FROM
THE RECORDS OF THE UNITED STATES PATENT AND TRADEMARK
OFFICE OF THOSE PAPERS OF THE BELOW IDENTIFIED PATENT
APPLICATION THAT MET THE REQUIREMENTS TO BE GRANTED A
FILING DATE UNDER 35 USC 111.**

APPLICATION NUMBER: 60/637,370

FILING DATE: December 17, 2004

**PRIORITY DOCUMENT
SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH
RULE 17.1(a) OR (b)**

**By Authority of the
COMMISSIONER OF PATENTS AND TRADEMARKS**



**P. SWAIN
Certifying Officer**

Please type a plus sign (+) inside this box → +

PTO/SB/16 (02-01)

Approved for use through 10/31/2002. OMB 0651-0032
Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PROVISIONAL APPLICATION FOR PATENT COVER SHEET

This is a request for filing a PROVISIONAL APPLICATION FOR PATENT under 37 CFR 1.53 (c).

Express Mail Label No. EV 507 975 663 US Date of Deposit: 17 December 2004

17513 U.S. PTO
60/637370

121704

INVENTOR(S)		
Given Name (first and middle [if any]) NEVENKA JUN	Family Name or Surname DIMITROVA FAN	Residence YORKTOWN HEIGHTS, NY SHENZHEN, GUANGDONG PROV. CN

Additional inventors are being named on the _____ separately numbered sheets attached hereto

TITLE OF THE INVENTION (280 characters max)

CONTINUOUS FACE RECOGNITION WITH ONLINE LEARNING

CORRESPONDENCE ADDRESS

Direct all correspondence to:

Customer Number 24737

OR Type Customer Number here

<input checked="" type="checkbox"/> Firm or Individual Name	PHILIPS INTELLECTUAL PROPERTY & STANDARDS				
Address	345 SCARBOROUGH ROAD				
Address	P. O. Box 3001				
City	BRIARCLIFF MANOR	State	NY	ZIP	10510
Country	USA	Telephone	(914) 333-9611	Fax	(914) 332-0615

24737

ENCLOSED APPLICATION PARTS (check all that apply)

- Specification Number of Pages 28 CD(s), Number _____
 Drawing(s) Number of Sheets 6 Other (specify) _____
 Application Data Sheet. See 37 CFR 1.76


METHOD OF PAYMENT OF FILING FEES FOR THIS PROVISIONAL APPLICATION FOR PATENT (check one)

- Applicant claims small entity status. See 37 CFR 1.27.
 A check or money order is enclosed to cover the filing fees
 The Commissioner is hereby authorized to charge filing fees or credit any overpayment to Deposit Account Number: 14-1270 160.00
 Payment by credit card. Form PTO-2038 is attached.

FILING FEE AMOUNT (\$)

The invention was made by an agency of the United States Government or under a contract with an agency of the United States Government.

- No.
 Yes, the name of the U.S. Government agency and the Government contract number are: _____

Respectfully submitted,
SIGNATURE 

Date 12/17/2004

REGISTRATION NO.: 28,613

(If appropriate)

Docket Number: US040102

TYPED or PRINTED NAME EDWARD W. GOODMAN

TELEPHONE (914) 333-9611

USE ONLY FOR FILING A PROVISIONAL APPLICATION FOR PATENT

This collection of information is required by 37 CFR 1.51. The information is used by the public to file (and by the PTO to process) a provisional application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 6 hours to complete, including gathering, preparing, and submitting the complete provisional application to the PTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, Washington, D.C. 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Box Provisional Application, Assistant Commissioner for Patents, Washington, D.C. 20231.

CONTINUOUS FACE RECOGNITION WITH ONLINE LEARNING

This application claims priority to U.S. provisional patent application 60/541,206, entitled "Continuous Face Recognition With Online Learning" of Nevenka Dimitrova and Jun Fan, filed February 2, 2004.

The contents of the above-identified U.S. provisional patent application 60/541,206, entitled "Continuous Face Recognition With Online Learning" of Nevenka Dimitrova and Jun Fan, filed February 2, 2004, are hereby incorporated by reference herein.

10 The invention generally relates to face recognition. More particularly, the invention relates to improvements in face recognition, including online learning of new faces.

15 Face recognition has been an active area of research, with many techniques currently available. One such technique uses a probabilistic neural network (generally "PNN") to determine whether it recognizes an input vector representing a face detected in a video stream or other image. The PNN determines whether a face is "known" or "unknown" by comparison of the input vector with a fixed number of known faces with which the PNN has been trained. If a comparison results in a sufficiently high confidence value, for example, the face is deemed to be that of the corresponding face in the database.

20 If the comparison does not, the input face is simply deemed to be "unknown" and discarded. PNNs are generally described, for example, in "Probabilistic Neural Network for Pattern Classification", by P. K. Patra et al., Proceedings of the 2002 International Joint Conference on Neural Networks (IEEE IJCNN '02), May 2002, Vol. II, pp. 1200-1205, the contents of which are hereby incorporated by reference herein.

25 One difficulty in prior techniques applying PNN to face recognition is that input
faces are only compared to faces in the pre-trained database. In other words, a face can
only be determined to be "known" if it is found to correspond to one of the faces used to
train the PNN. Thus, the same input face may be repeatedly determined to be "unknown"
if it is not in the database, even though the same face has previously been detected by the
30 system.

U.S. Patent Application Publication 2002/0136433 A1 ("433 publication") describes a face recognition system that applies online training for unknown faces in an

"adaptive eigenface" system. According to the '433 publication, an unknown face that is detected is added to the class of known faces. The '433 publication also refers to tracking the face so that multiple images of the unknown face may be added to the database. However, the '433 publication does not teach selectivity in determining whether or not to add unknown faces to the database. Thus, the '433 database may rapidly expand with new faces and also slow down the performance of the system. While capture of all unknown images may be desirable for certain applications (such as surveillance, where it may be desirable to capture every face for later recognition), it may be undesirable in others. For example, in a video system where rapid identification of prominent faces is important, indiscriminate expansion of the database may be undesirable.

The present invention includes, among other things, addition of new faces to a database or the like used in face recognition and keeps learning new faces. When a new face is added to the database, it may be detected as a "known" face when it is found again in the input video subsequently received. One aspect discriminates which new faces are added to the database by applying rules to ensure that only new faces that persist in the video are added to the database. This eliminates "spurious" or "fleeting" faces from being added to the database.

A side note is made here regarding terminology as utilized in the description below: In general, a face is considered "known" by a system if data regarding the facial features is stored in the system. In general, where a face is "known", an input containing the face may be recognized by the system as corresponding to the stored face. For example, in a PNN based system, a face is "known" if there is a category corresponding to the face and is considered "unknown" if there is no such category. (Of course, the existence of a category corresponding to a face does not necessarily mean that the processing will always determine a match or hit, since there may be "misses" between an input known face and its category.) A "known" face will generally be given an identifier by the system, such as a generic label or reference number. (As will be seen labels F1, F2, ..., FN in Figs. 2 and 6 and FA in Fig 6 represent such generic identifiers in the system.) A system may have stored data regarding facial features and such system identifiers or labels for the faces without necessarily having the identity of the person (such as the person's name). Thus, a system may "know" a face in the sense that it includes stored facial data for the face without necessarily having data relating to the personal identification of the face. Of

course, a system may both "know" a face and also have corresponding personal identification data for the face.

Thus, the invention comprises a system having a face classifier that provides a determination of whether or not a face image detected in a video input corresponds to a known face in the classifier. The system adds an unknown detected face to the classifier when the unknown detected face persists in the video input in accordance with one or more persistence criteria. The unknown face thus becomes known to the system.

The face classifier may be, for example, a probabilistic neural network (PNN), and the face image detected in the video input is a known face if it corresponds to a category in the PNN. When the persistence criteria is met for an unknown face, the system may add 10 the unknown face to the PNN by addition of a category and one or more pattern nodes for the unknown face to the PNN, thereby rendering the unknown face to be known to the system. The one or more persistence criteria may comprise detection of the same unknown face in the video input for a minimum period of time.

15 The invention also comprises a like method of face classification. For example, a method of face recognition comprising the steps of: determining whether or not a face image detected in a video input corresponds to a known face in storage, and adding an unknown detected face in storage when the unknown detected face persists in the video input in accordance with one or more persistence criteria.

20 The invention also comprises like techniques of face classification using discrete images, such as photos. It also provides for adding an unknown face (in either the video or discrete image case) when a face in at least one image meets one or more prominence criteria, e.g., a threshold size.

25 The preferred exemplary embodiment of the present invention will hereinafter be described in conjunction with the appended drawings, where like designations denote like elements, and:

Fig. 1 is a representative block diagram of a system according to an embodiment of the invention;

30 Fig. 1a is a representative diagram of a different level of the system of Fig. 1;
Fig. 2 is an initially trained modified PNN of a component of the system of Fig. 1;
Fig. 3 is a more detailed representation of a number of components of the system of

Fig. 1;

Fig. 3a is a vector quantization histogram created for a face image in accordance with a feature extraction component as in Fig. 3;

Fig. 4 is a representative one-dimensional example used in showing certain results based on a probability distribution function;

5 Fig. 5 shows a modification of the example of Fig. 4; and

Fig. 6 is the modified PNN of Fig. 2 including a new category created by online training.

As noted above, the present invention comprises, among other things, face recognition that provides for online training of new (i.e., unknown) faces that persist in a 10 video image. The persistence of a new face in a video image is measured by one or more factors that provide, for example, confirmation that the face is a new face and also provides a threshold that the face is one sufficiently significant to warrant addition to the database for future determinations (i.e., become a "known" face).

15 Fig. 1 depicts an exemplary embodiment of the invention. Fig. 1 is representative of both a system and method embodiment of the invention. The system terminology will be used below to describe the embodiment, although it is noted that the processing steps described below also serve to describe and illustrate the corresponding method embodiment. As will be readily apparent from the description below, video inputs 20 and the sample face images 70 above the top dotted line (portion A) are inputs to the system 10, which may be stored in a memory of system 10 after receipt. Processing blocks inside the 20 dotted lines (portion "B") comprise processing algorithms that are executed by system 10 as described further below.

25 As will be readily appreciated by those of skill in the art, the processing algorithms of system 10 in portion B may reside in software that is executed by one or more processors and which may be modified by the system over time (e.g., to reflect the online training of the MPNN described below). As will also become clear from the description below, the inputs to various processing block algorithms are provided by the output of other processing blocks, either directly or through an associated memory. (Fig. 1a provides a simple representative embodiment of the hardware and software components that support 30 the processing of system 10 represented in Fig. 1. Thus, the processing of system 10 represented by the blocks in portion B of Fig. 1 may be performed by the processor 10a in conjunction with associated memory 10b and software 10c in Fig. 1a.)

Scanned with CamScanner

Scanned with CamScanner on 01/27/2005

histogram technique of feature extraction is well known in the art and also described further below in the context of analogous feature extraction in block 35 for input video images. Thus, input feature vector X for each sample image will have a number of dimensions determined by the vector codebook used (33 in the particular example below).

5 After input feature vector X of a sample image is extracted, it is normalized by classifier trainer 80. Classifier trainer 80 also assigns the normalized X as a weight vector W to a separate pattern node in the MPNN 42. Thus, each pattern node also corresponds to a sample image of one of the faces. Trainer 80 connects each pattern node to a node created for the corresponding face in the category layer. Once all sample input images are
10 received and processed in like manner, the MPNN 42 is initially trained. Each face category will be connected to a number of pattern nodes, each pattern node having a weight vector corresponding to a feature vector extracted from a sample face image for the category. Collectively the weight vectors of the pattern nodes for each face (or category) create an underlying probability distribution function (PDF) for the category.

15 Fig. 2 is a representation of an MPNN 42 of face classifier 40 as initially offline trained 90 by the classifier trainer 80. A number n_1 of the input sample images output by block 70 correspond to face F1. Weight vector W_{11} assigned to first pattern node equals a normalized input feature vector extracted from first sample image of F1; weight vector W_{12} assigned to second pattern node equals a normalized input feature vector extracted
20 from second sample image of F1; ...; and weight vector W_{1n_1} assigned to n_1^{th} pattern node equals a normalized input feature vector extracted from n_1^{th} sample image of F1. The first n_1 pattern nodes are connected to the corresponding category node F1. Similarly, a number n_2 of the input sample images correspond to face F2. The next n_2 pattern nodes having weight vectors $W_{21} - W_{2n_2}$, respectively, are created in like manner
25 using the n_2 sample images of F2. The pattern nodes for face F2 are connected to category F2. Subsequent pattern nodes and category nodes are created for subsequent face categories in like manner. In Fig. 2, the training uses multiple sample images for N different faces.

An algorithm for creating the initially trained MPNN of Fig. 2 is now briefly
30 described. As noted above, for a current sample face image input at block 70, feature extractor 75 first creates a corresponding input feature vector X (which in the particular embodiment is a VQ histogram, described below). Classifier trainer 80 converts this input

feature vector to a weight vector for a pattern node by first normalizing the input feature vector by dividing the vector by its respective magnitude:

$$X' = X \cdot \left(1 / \sqrt{\sum X^2} \right) \quad (1)$$

- The current sample image (and thus currently corresponding normalized feature vector X') corresponds to a known face F_j , where F_j is one of the faces F_1, F_2, \dots, F_N of the training. Also, as noted, there will generally be a number of sample images for each known face in the stream of sample faces of block 70. Thus, current sample image will generally be the m -th sample image corresponding to F_j output by block 70. The normalized input feature vector X' is thus assigned as a weight vector to the m -th pattern node for category F_j :

$$W_{jm} = X' \quad (2)$$

- The pattern node with weight vector W_{jm} is connected to the respective category node F_j . The other sample face images input by block 70 are converted to input feature vectors in feature extraction block 75 and processed in like manner by classifier trainer 80 to create the initially configured MPNN 42 of face classifier shown in Fig. 2.

- For example, referring back to Fig. 2, if the current sample image input by block 70 is a first sample image for face F_1 , then feature extractor 75 creates input feature vector X for the image. Classifier trainer 80 normalizes input feature vector and assigns it as the weight vector W_{11} for the first pattern node for F_1 . The next sample image may be for third sample image for face F_9 . After extraction of an input feature vector X for this next sample image at block 75, classifier trainer 80 normalizes the feature vector and then assigns the normalized feature vector as weight vector W_{93} for the third pattern node for F_9 (not shown). Some input images later, another sample image in the training may again be for F_1 . This image is processed in like manner and assigned as weight vector W_{12} for the second pattern node for F_1 .

- All sample face images 70 are processed in like manner, resulting in the initially trained MPNN 42 of classifier 40 of Fig. 2. After such initial offline training 90, face classifier 40 comprises an MPNN 42 having pattern layer and category layer resulting from offline training and reflecting the faces used in the offline training. Such faces comprise the initially "known" faces of the offline trained MPNN-based system.

As described further below, input nodes I_1, I_2, \dots, I_M will receive a feature vector of a detected face image and determine if it corresponds to a known face category. Thus,

each input node is connected to each pattern node and the number of input nodes equals the number of dimensions in the feature vectors (33 in the particular example below).

The training of MPNN may be done as a sequence of input sample images, as described above, or multiple images may be processed simultaneously. Also, it is clear from the above description that the order of input of the sample face images is irrelevant. Since the face category is known for each sample image, all samples for each known face may be submitted in sequence, or they may be processed out of order (as in the example given above). In either case, the final trained MPNN 42 will be as shown in Fig. 2.

It is noted that the MPNN as configured immediately after such initial offline training of system 10 is analogous to those in prior art PNN systems that only use offline training. For example, such offline training 90 may be done in accordance with the above-cited document by Patra et al.

It is noted here (and further described below) that the present invention does not necessarily require offline training 90. Instead the MPNN 42 may be built up using solely online training 110, also further described below. However, for the currently described embodiment, the MPNN 42 is first trained using offline 90 training and is as shown in Fig. 2. After the initial offline training 90 of MPNN 42 as described above, the system 10 is used to detect a face in a video input 20 and, if detected, to determine whether the detected face corresponds to a known face of one of the categories of the MPNN 42. Referring back to Fig. 1, video input 20 is first subject to an existing technique of face detection 30 processing, which detects the presence and location of a face (or faces) in the video input 20. (Thus, face detection processing 30 merely recognizes that an image of a face is present in the video input, not whether it is known.) System 10 may use any existing technique of face detection.

Face detection algorithm 30 may thus utilize the known application of AdaBoost to rapid object detection as described in "Rapid Object Detection Using A Boosted Cascade of Simple Features" by P. Viola and M. Jones, Proceedings of the 2001 IEEE Conference on Computer Vision and Pattern Recognition (IEEE CVPR '01), Vol. I, pp. 511-518, Dec. 2001, the contents of which are hereby incorporated by reference herein. The basic face detection algorithm 30 used may be as described in Viola, namely, it is structured in cascaded stages, with each stage being a strong classifier and each stage comprised of several weak classifiers, each weak classifier corresponding to a feature of the image.

Input video images 20 are scanned from left to right, top to bottom, and rectangles of different sizes in the image are analyzed to determine whether or not it contains a face. Thus, stages of the classifier are applied in succession to a rectangle. Each stage yields a score for the rectangle, which is the sum of the responses of the weak classifiers comprising the stage. (As noted below, scoring for the rectangle typically involves looking into two or more sub-rectangles.) If the sum exceeds a threshold for the stage, the rectangle proceeds to the next stage. If the rectangle's scores pass the thresholds for all stages, it is determined to include a face portion, and the face image is passed to feature extraction 35. If the rectangle is below the threshold for any stage, the rectangle is discarded and the algorithm proceeds to another rectangle in the image.

The classifier may be constructed as in Viola by adding one weak classifier at a time that are evaluated using a validation set to build up the stages or strong classifiers. The newest weak classifier is added to the current stage under construction. Each round t of boosting adds a rectangular feature classifier h_t to the current set of features in the strong classifier under construction by minimizing:

$$E_t = \sum_i D_t(i) \exp(-\alpha_t y_i h_t(x_i)) \quad (3)$$

The above equation 3 is equivalent to the one used in Viola's procedure, and E_t represents a weighted error associated with the t^{th} rectangular feature classifier h_t being evaluated using rectangular training example x_i . (The lower case notation "x_i" used for the rectangular example distinguishes it from the feature vector notation X of images used in the MPNN.) Fundamentally $h_t(x_i)$ is a weighted sum of sums of pixels in particular rectangular sub-regions of training example x_i . If $h_t(x_i)$ exceeds a set threshold, then the output of $h_t(x_i)$ for example x_i is 1 and, if not, the output of $h_t(x_i)$ is -1. Because h is restricted in the above equation to +1 or -1, the variable α_t is the influence (magnitude) of this weak hypothesis h on the strong classifier under construction. Also, $y_i \in [-1, 1]$ is the target label of example x_i (that is, whether x_i is a negative or positive example of feature h, which is objectively known for the examples of the training set). D is a weighting factor for the ith example for the h_t feature.

Once the minimum E is determined in this manner, the corresponding rectangular feature classifier h (as well as its magnitude α) is used to construct the new weak classifier.

A custom decision threshold for h is also determined using the training set and based on the distribution of positive and negative examples. The threshold is selected that best partitions the positive and negative examples based on design parameters. (The threshold is referred to in the above-referenced Viola document as θ_j .) As noted, the weak classifier is also comprised of α , which is a real-valued number that denotes how much influence the rectangular feature classifier h selected has on the strong classifier under construction (and is determined from the error E determined in the training). When implemented, an input rectangular portion of an image is also typically analyzed by h based on the weighted sum of pixels in two or more sub-rectangles of the input rectangle, and the output of h is set to 1 if the threshold (as determined from the training) is exceeded for the input rectangle and $h=-1$ if it does not. The output of new weak classifier is the binary output of h times the influence value α . The strong classifier is comprised of the sum of the weak classifiers added during the training.

Once a new weak classifier is added, if the classifier's performance (in terms of detection rates and false alarm rates) meets the desired design parameters for the validation set, then the newly added weak classifier completes the stage under construction, since it adequately detects its respective feature. If not, another weak classifier is added and evaluated. Once stages are constructed for all desired features and perform in accordance with the design parameters for the validation set, the classifier is completed.

A modification of the above-described structure of the Viola weak classifiers may alternatively be utilized for face detector 30. In the modification, α is folded into h during the selection of h for the new weak classifier. The new weak classifier h (which now incorporates α) is selected by minimizing E in manner analogous to that described above. As to the implementation of the weak classifier, "boosting stumps" are utilized in the modification. Boosting stumps are decision trees that output the left or right leaf value based on the decision made at the non-leaf parent node. Thus, weak classifier is comprised of a decision tree that outputs one of two real values (one of two leafs c_{left} and c_{right}) instead of 1 and -1. Weak classifier is also comprised of a custom decision threshold, described below. For an input rectangle portion of an image, the selected rectangular feature classifier h is used to determine if the weighted sum of the sums of pixel intensities between sub-rectangular regions of the input rectangle is greater than the threshold. If greater, c_{left} is output from the weak classifier, if less, c_{right} is output.

Leaves c_{left} and c_{right} are determined during the training of the selected h , based on how many positive and negative examples are assigned to the left and right partitions for a given threshold. (Examples are objectively known to be positive or negative because ground truth on the training set is known.) The weighted sum of sums from the rectangles are evaluated over the entire sample set, thus giving a distribution of difference values, which is then sorted. From the sorted distribution and in view of the required detection and false alarm rates, the goal is to select a partition wherein most positive examples fall to one side and most negative examples fall to the other. For the sorted distribution, the optimum split (giving the custom decision threshold used for the weak classifier) is done by choosing a partition that minimizes T in the following equation:

$$T = 2 \left(\sqrt{W_+^{\text{Left}} W_-^{\text{Left}}} + \sqrt{W_+^{\text{Right}} W_-^{\text{Right}}} \right) \quad (4)$$

where W denotes the weight of the examples in the training set that fall to the left or right of the partition under consideration that are either "positive" or "negative".

The selected partition (that minimizes T) creates the custom decision threshold; also, c_{left} and c_{right} are computed from the training data distribution according to the equations:

$$c_{\text{left}} = \frac{1}{2} \ln \left(\frac{W_+^{\text{Left}} + \epsilon}{W_-^{\text{Left}} + \epsilon} \right) \quad \text{and} \quad c_{\text{right}} = \frac{1}{2} \ln \left(\frac{W_+^{\text{Right}} + \epsilon}{W_-^{\text{Right}} + \epsilon} \right) \quad (5, 6)$$

where W now denotes the weight of the examples that are assigned to the left or right of the selected partition that are either "positive" or "negative" (and ϵ is a smoothing term to avoid numerical problems caused by large predictions). These values serve to keep the weights of the next iteration of weak classifier balanced, that is, keep the relative weights of positive and negative examples on each side of the boundary substantially equal.

As noted, although weak classifiers may be structured as in Viola, alternatively they may be structured as decisions stumps described directly above. In addition, it is noted that the training of either weak classifier may use alternative techniques. According to one technique, to test the weak classifier currently being added, the examples of the validation set are scanned through all previously added weak classifiers of prior stages and weak classifiers previously added to the current stage. However, once a prior weak classifier is adopted and scored, the score does not change. Thus, in a more efficient alternative technique, the rectangles that pass through all prior stages and their scores for the prior stages are stored. Rather than running the examples through all prior stages, the prior



Scanned by L. M. Nisoto from the IFW Image Database on 01/27/2005

X_D is input to MPNN 42 via the input layer nodes and MPNN 42 evaluates its correspondence with each face category using the weight vectors in the pattern nodes. MPNN 42 compares X_D and a known face category (F_1, F_2, \dots) by determining a separate PDF value for each category. First, the input layer normalizes the input vector X_D , (by dividing it by its magnitude) so that it is scaled to correspond with prior normalization of the weight vectors of the pattern layer during offline training:

$$X'_D = X_D \cdot \left(1 / \sqrt{\sum X_D^2}\right) \quad (7)$$

Second, in the pattern layer, MPNN 42 performs a dot product between the normalized input vector X'_D and the weight vector W of each pattern node shown in Fig. 2, thus resulting in an output vector value Z for each pattern node:

$$Z_{11} = X'_D \cdot W_{11}, \quad (8a)$$

$$Z_{12} = X'_D \cdot W_{12}, \quad (8b)$$

$$Z_{N_{n,N}} = X'_D \cdot W_{N_{n,N}} \quad (8n)$$

where the reference notations for the weight vectors W for the pattern nodes (and thus the resulting output vectors Z) are as shown in Fig. 2 and as described above with respect to the offline training.

Finally, the output values of pattern nodes corresponding to each category are aggregated and normalized to determine a value of the PDF (function f) for input vector X_D for each respective category. Thus, for the j -th category F_j , output values $Z_{j1} - Z_{jn,j}$ for pattern nodes of the j -th category are used, where n_j is the number of pattern nodes for category j . The PDF value f is calculated for a category F_j under consideration as follows:

$$f_{Fj}(X_D) = \sum_{l=1}^{n_j} (\exp[(Z_{jl}-1)/\sigma^2])/n_j \quad (9)$$

where σ is the smooth factor. Using equation 9 for $j=1$ to N , PDF values $f_{F1}(X_D), \dots, f_{FN}(X_D)$ are calculated for categories F_1, \dots, F_N , respectively, using the output values Z of the pattern nodes corresponding to each respective category. Because the PDF value f for each category is based on a sum of the output values Z of the category, it follows that the greater the value f for a category, the greater the correspondence between X_D and the weight vectors for that category.

The MPNN 42 then selects the category (designated the i th category or F_i) that has the largest value f for input vector X_D . Selection of the i th category by the MPNN 42 uses one of the implementations of the Bayes Strategy, which seeks the minimum risk cost based on the PDF. Formally, the Bayes decision rule is written as:

5 $d(X_D) = F_i \text{ if } f_{F_i}(X_D) > f_{F_j}(X_D) \quad \forall i \neq j \quad (10)$

Category F_i having the largest PDF (as measured by f) for input vector X_D provides a determination that input vector X_D (corresponding to face segment 42a) potentially matches known face category F_i . Before actually deeming there is a match, the MPNN 42 generates a confidence measurement, which compares the PDF of vector X_D for 10 the potential matching category i with the sum of the PDFs of vector X_D for all categories:

$$C_i = f_{F_i}(X_D) / \left(\sum_{j=1}^N f_{F_j}(X_D) \right) \quad (11)$$

If the confidence measurement surpasses a confidence threshold (e.g., 80%), then a match between input vector X_D and category i is found by the system. Otherwise it is not.

However, the confidence measurement based on the decision function result as 15 described directly above can result in undesirably high confidence measurements in cases where the largest PDF value f for an input vector is nonetheless too low for a match with the category to be declared. This is due to the confidence measurements as calculated above being generated by comparing the relative results from the PDF output of the categories for a given input vector. A simple generic example in one-dimension illustrates 20 this:

Fig. 4 represents the PDF of two categories (Cat1, Cat2). The PDF function for each category is generically represented in Fig. 4 as “ $p(X|Cat)$ ” (or the probability that input feature vector X belongs to category Cat) versus the one dimensional feature vector X . Three separate one dimensional input feature vectors X_{Ex1} , X_{Ex2} , X_{Ex3} are shown which 25 are used to illustrate how undesirably high confidence values may result. For input vector X_{Ex1} , the largest PDF value corresponds to category Cat1 (i.e., $p(X_{Ex1}|Cat1) \approx 0.1$ and $p(X_{Ex1}|Cat2) \approx 0.02$). By applying a Bayes rule analogous to that given in equation 10, Cat1 is thus selected. Also, a confidence measurement may be calculated for Cat1 for X_{Ex1} analogous to that given in equation 11:

30 $\text{Conf}_i_{Ex1} = p(X_{Ex1}|Cat1) / [p(X_{Ex1}|Cat1) + p(X_{Ex1}|Cat2)] \quad (12)$
 $\approx 0.1 / [0.1 + 0.02] = 83\%$

However, since the PDF values for input feature vector X_{Ex1} are very low (0.1 for Cat1 and lower for Cat2), this implies that the correspondence between the input vector and the weight vectors in the pattern nodes is small, and that X_{Ex1} should therefore be identified as an "unknown" category.

5 Other like undesirable results are also evident from Fig. 4: Referring to input feature vector X_{Ex2} , it is clearly appropriate to match it with category Cat1 since it corresponds to the maximum value of Cat1. Also, calculation of the confidence value Confi_Ex2 in manner analogous to equation 12 results in a confidence measurement of approximately 66%. However, Confi_Ex2 should not be lower than Confi_Ex1, since X_{Ex2} is much closer than X_{Ex1} to the maximum value of the PDF for Cat1. Another 10 undesirable result is shown for X_{Ex3} , where Cat2 is selected with a confidence value of approximately 80%, even though X_{Ex3} is likewise far to one side of the maximum value of the PDF for Cat2.

15 Fig. 5 exemplifies a technique for avoiding such undesirable outcomes when treating low PDF values for a given input feature vector. In Fig. 5, a threshold is applied to each of the categories Cat1, Cat2 of Fig. 4. In addition to choosing the category having the largest PDF value, an input feature vector X must meet or exceed the threshold for the category before it is deemed a match. The threshold may be different for each category. For example, the threshold may be a certain percentage of the maximum value of the PDF 20 for the category (e.g., 70%).

As seen in Fig. 5, Cat1 is again the category having the largest PDF value for feature vector X_{Ex1} . However, $p(X_{Ex1}|Cat1) \approx 0.1$, and does not surpass the threshold for Cat1, which is approximately 0.28. Thus, feature vector X_{Ex1} is determined to be "unknown". Likewise, since the PDF value of X_{Ex3} does not surpass the threshold for 25 Cat2, X_{Ex3} is determined to be "unknown". However, since the PDF value for X_{Ex2} surpasses the threshold for Cat1, Cat1 is selected for X_{Ex2} , with a confidence level of 66% as calculated above.

It is clear that analogous undesirable scenarios can arise when in the case of multi-dimensional cases (such as the 33 dimensional case in the exemplary embodiment). For 30 example, the PDF value for the largest category for an input multi-dimensional feature vector may nonetheless be too low to declare a category match. However, when the largest

value does not surpass the confidence threshold, then the input vector is again deemed unknown.

Processing the determination of whether the face is known or unknown is separately shown as processing determination 50 in Fig. 1. Block 50 may include the 5 modified Bayes decision rule (Equations 13 and 14) and the subsequent confidence determination (Equation 11) as described immediately. However, although block 50 is shown separately from face classifier 40 for conceptual clarity, it is understood that the Bayes decision algorithm and confidence determination is typically part of face classifier 10 40. This decision processing may be considered part of the MPNN 42, although it may alternatively be considered a separate component of face classifier 40.

If the face image is determined by determination 50 to be unknown, Fig. 1 shows that the face is not simply discarded but rather the processing turns to a persistence 15 decision block 100. As described in more detail below, the video input 20 having the unknown face is monitored using one or more criteria to determine if the same face persists or is otherwise prevalent in the video. If it does, then the feature vectors X_D for one or 20 more face images of the unknown face received via input 20 are sent to the trainer 80. Trainer 80 uses the data for the face images to train the MPNN 42 in face classifier 40 to include a new category for the face. Such "online" training of the MPNN 42 ensures that a prominent new (unknown) face in the video will be added as a category in the face 25 classifier. Thus the same face in subsequent video inputs 20 may be detected as a "known" face (i.e., corresponding to a category, although not necessarily "identified" by name, for example).

As noted, when the face is determined to be unknown in block 50, persistence 25 processing 100 is initiated. Video input 20 is monitored to determine if one or more conditions are satisfied, indicating that the MPNN 42 will be online trained using images of the unknown face. The one or more conditions may indicate, for example, that the same unknown face is continuously present in the video for a period of time. Thus, in one embodiment of the persistence processing 100, the unknown face detected is tracked in the video input using any well-known tracking technique. If the face is tracked in the video 30 input for a minimum number of seconds (e.g., 10 seconds), then the face is deemed to be persistent by processing block 100 ("yes" arrow).

- In the case where the sequence lasts for a period of time that is continuous, the processing is straightforward. In that case, some or all of the feature vectors X_D for the face segments of video input 20 may be stored in a buffer memory and, if the minimum period of time is exceeded, used in online training as described further below. In other cases, for example, a face may appear for very short periods of time in non-consecutive video segments, but which aggregate to exceed the minimum period of time. (For example, where there are rapid cuts between actors engaged in a conversation.) In that case, multiple buffers in persistence block 100 may each store feature vectors for unknown face images for a particular unknown face, as determined by above conditions 1-3.
- 5 Subsequent face images that are determined to be "unknown" by MPNN are stored in the appropriate buffer for that face, as determined by criteria 1-3. (If an unknown face does not correspond to those found in an existing buffer, it is stored in a new buffer.) If and when a buffer for a particular unknown face accumulates enough feature vectors for face images over time to exceed the minimum period of time, the persistence block 100 releases 10 the feature vectors to classifier trainer 80 for online training 110 for the face in the buffer.
- 15

If the sequence of faces for an unknown face is determined not to meet the persistence criteria (or a single persistence criterion), then the processing of the sequence is terminated and any stored feature vectors and data relating to the unknown face are discarded from memory (processing 120). In the case where image segments are 20 accumulated for different faces over time in different buffers as described above, the data in any one buffer may be discarded if, after a longer period of time (e.g., 5 minutes), the face images accumulated over time does not exceed the minimum period.

If a face in the video input determined to be unknown satisfies the persistence processing, then system 10 performs an online training 110 of the MPNN 42 to include a 25 category for the unknown face. For convenience, the ensuing description will focus on online training for unknown face "A" that satisfies persistence block 100. As described above, in the determination of the persistence of face A, the system stores a number of feature vectors X_D for images of face A from the sequence of images received via video input 20. The number of feature vectors may be for all of the faces of A in the sequence 30 used in the persistence determination, or a sample. For example, input vectors for 10 images in the sequence of face A may be utilized in the training.

For a persistent face A, system processing returns to training processing 80 and, in this case, online training 110 of MPNN 42 of face classifier 40 to include face A. The 10 feature vectors used (for example) in the online training for face A may be those having the lowest variance from all the input vectors for the images in the sequence, that is, the 10
5 input vectors having closest to the average in the buffer. Online training algorithm 110 of trainer 80 trains the MPNN 42 to include a new category FA for face A having pattern nodes for each of the images.

The online training of new category FA proceeds in analogous manner for the initial offline training of the MPNN 42 using sample face images 70. As noted, the feature
10 vectors X_D for the images of face A are already extracted in block 35. Thus, in the same manner as the offline training, classifier trainer 80 normalizes the feature vectors of FA and assigns each one as a weight vector W of a new pattern node for category FA in the MPNN. The new pattern nodes are connected to a category node for FA.

Fig. 6 shows the MPNN of Fig. 2 with new pattern nodes for new category FA.
15 The newly added nodes are in addition to the N categories and corresponding pattern nodes developed in the initial offline training using known faces discussed above. Thus, weight vector WA_1 assigned to first pattern node for F1 equals a normalized feature vector for a first image of FA received via video input 20; weight vector WA_2 assigned to second pattern node (not shown) for FA equals a normalized feature vector for a second sample
20 image of FA; ...; and weight vector WA_{n_A} assigned to n_A^{th} pattern node for FA equals a normalized feature vector for the n_1^{th} sample image of FA. By such online training, face A becomes a "known" face in MPNN. MPNN 42 is now capable of determining face A in a subsequent video input 20 is a "known" face using the detection and classification processing of Fig. 1 and described above. It is again noted that a face image A in a
25 subsequent video input 20 may be determined to be "known" in that it corresponds to a face category FA of MPNN. However, this does not necessarily mean that the face is "identified" in the sense that the name of face A is known to the system 10.

Other faces detected in the input video 20 and classified as "unknown" by system 10 in the manner described above are likewise processed by persistence processing 100. If
30 and when the one or more criteria applied in persistence block 100 is met by another face (e.g., face B), the trainer 80 online trains 110 the MPNN 42 in the manner described above for face A. After online training, MPNN 42 includes another category (with corresponding

While the invention has been described with reference to several embodiments, it will be understood by those skilled in the art that the invention is not limited to the specific forms shown and described. Thus, various changes in form and details may be made therein without departing from the spirit and scope of the invention as defined by the appended claims. For example, there are many alternative techniques that may be used in the present invention for face detection 30. An exemplary alternative technique of face detection as known in the art is further described in "Neural Network-Based Face Detection" by H.A. Rowley et al., IEEE Transactions On Pattern Analysis and Machine Intelligence", vol. 20, no. 1, pp. 23-38 (Jan., 1998).

In addition, other techniques of feature extraction may be used as alternatives to VQ histogram techniques described above. For example, the well-known "eigenface" technique may be used for comparing facial features. In addition, there are many variations of PNN classification that may be used as an alternative to the MPNN described above for face classification in which, for example, the online training techniques described above may be utilized. Also, there are many other techniques of face classification which may be used as alternatives to (or in techniques apart from) the MPNN technique utilized in the above exemplary embodiment, such as RBF, Naive Bayesian Classifier, and nearest neighbor classifier. The online training techniques, including the appropriate persistence and/or prominence criteria, may be readily adjusted to such alternative techniques.


Also, it is noted, for example, that the embodiment described above does not necessarily have to be initially offline trained with images of N different sample faces. The initial MPNN 42 may not have any offline trained nodes, and may be trained exclusively online with faces that meet the one or more persistence (or prominence) criteria, in the manner described above.

Also, persistence criteria other than those specifically discussed above fall within the scope of the invention. For example, the threshold time that a face needs to be present in a video input may be a function of video content, scene in the video, etc.

Thus, the particular techniques described above are by way of example only and not to limit the scope of the invention.

1/6

Fig. 1






Fig. 3

4/6

Fig. 1a





Fig. 3a



5/6




Fig. 4




Fig. 5




Fig. 6